

IA-64 Architecture Innovations

**Abbreviated Version of 2/23/99 IA-64
Architecture Disclosure**



Agenda

- Review content disclosed at IDF
 - ◆ Focus on benefits
 - ◆ Focus on “What’s New?”
- Branch Handling: Predication and Prediction
- Speculation
- Register Rotation & Loop Handling

Branch Handling

Traditional Arch



Dynamic Prediction

Static Prediction

Predication

IA-64



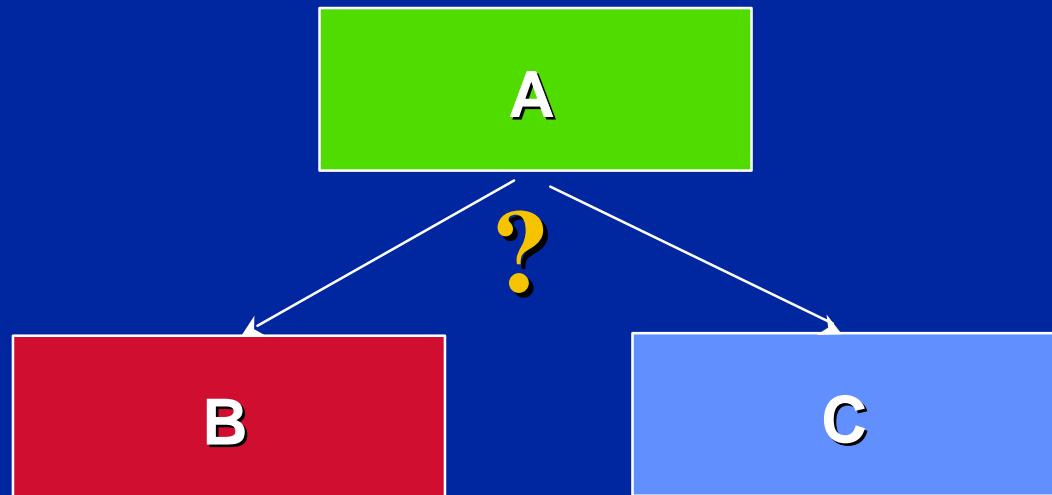
& Others

- Branches are forks in code
- Can indicate a decision
- Common in wide variety of applications



Review

Dynamic Branch Prediction

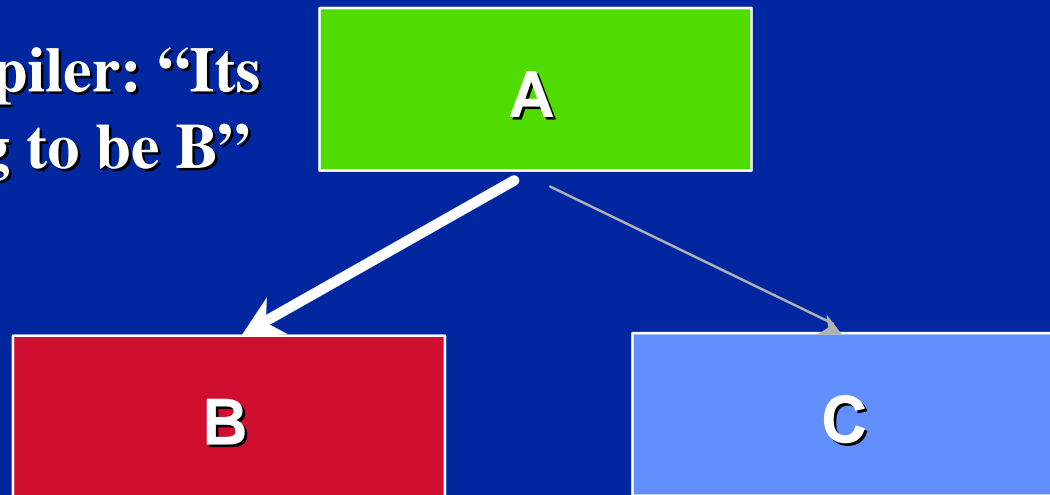


- Guesses either B or C
- Suffers performance penalty when mispredicted

New!

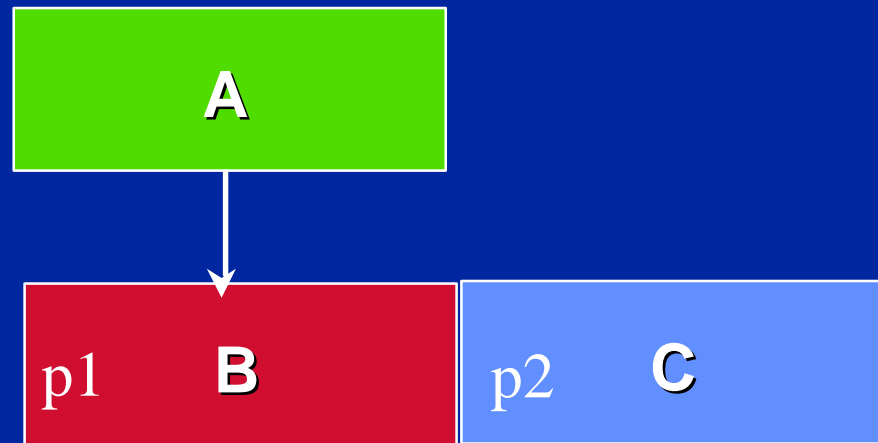
Static Branch Prediction

Compiler: “Its going to be B”



- Compiler knows which one is almost always taken or never taken
- Removes guesswork for processor, reduces mispredict penalties

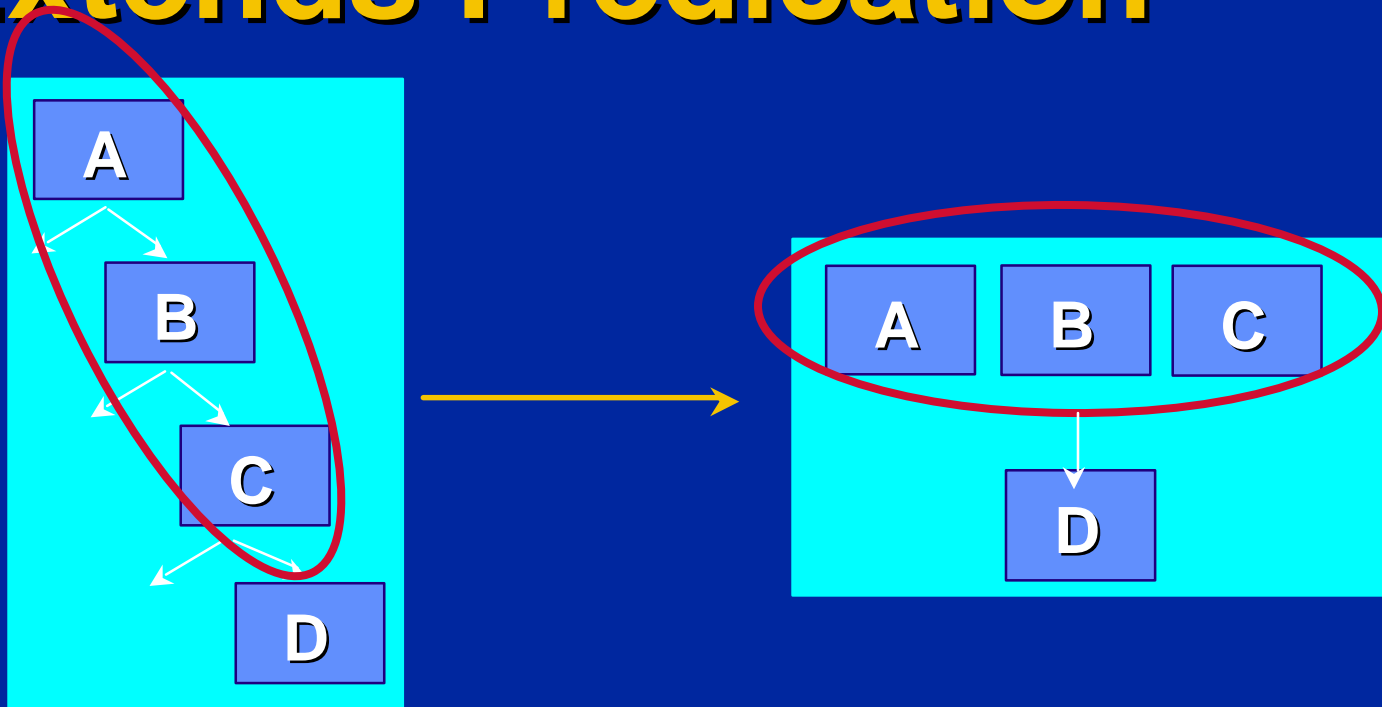
Predication



- Removes Branch, executes B&C in parallel
- Avoids possibility of mispredict
- Predicates (p1 & p2) are bits that turn on/off B & C
- Biggest benefit to code w/ hard to predict branches
 - Large server apps
 - Data sorting

New!

Parallel Compares: Extends Predication



Enables reduction from 7 to 5 cycles on queens loop

*Reduces critical path, further
increasing performance*

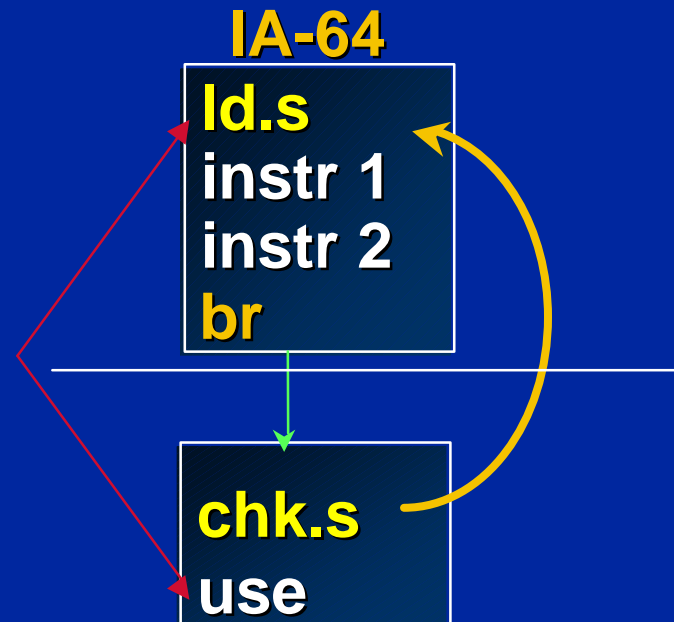
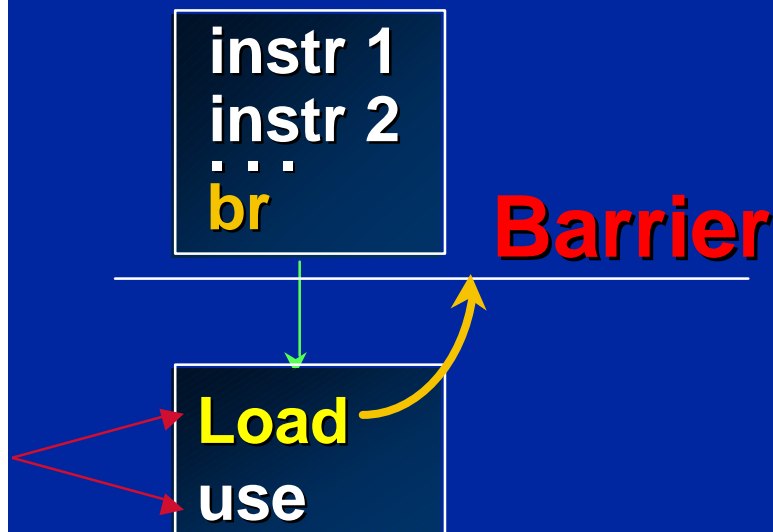
Predication Benefits

- Reduces branches and mispredict penalties
 - ◆ 50% fewer branches and 37% faster code*
- Parallel compares further reduce critical paths
- Greatly improves code with hard to predict branches
 - ◆ Large server apps- capacity limited
 - ◆ Sorting, data mining- large database apps
 - ◆ Data compression
- Traditional architectures' "bolt-on" approach can't efficiently approximate predication
 - ◆ Cmove: 39% more instructions, 30% lower performance*
 - ◆ Instructions must all be speculative

Review

Speculation Review

Traditional Architectures



**Allows elevation of load,
even above a branch**

- Moving things up enables them to be performed sooner
- Loads take time, need to separate from use

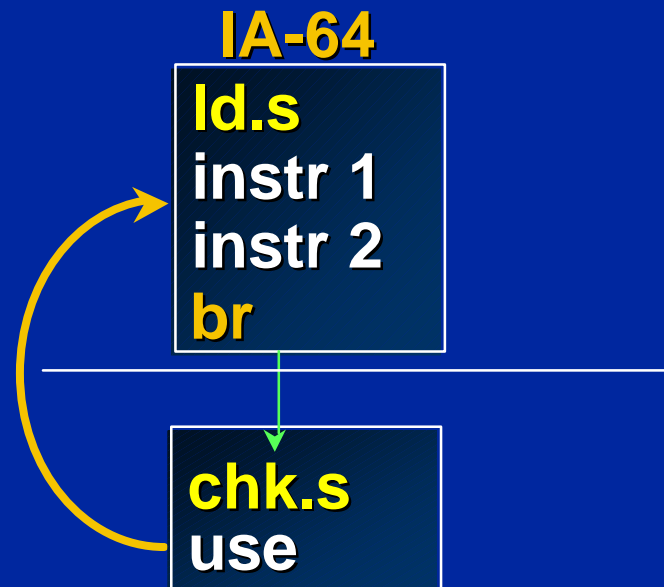
◆ Memory latency is a major problem in today's systems

◆ CPUs getting faster than memory



New!

Hoisting Uses



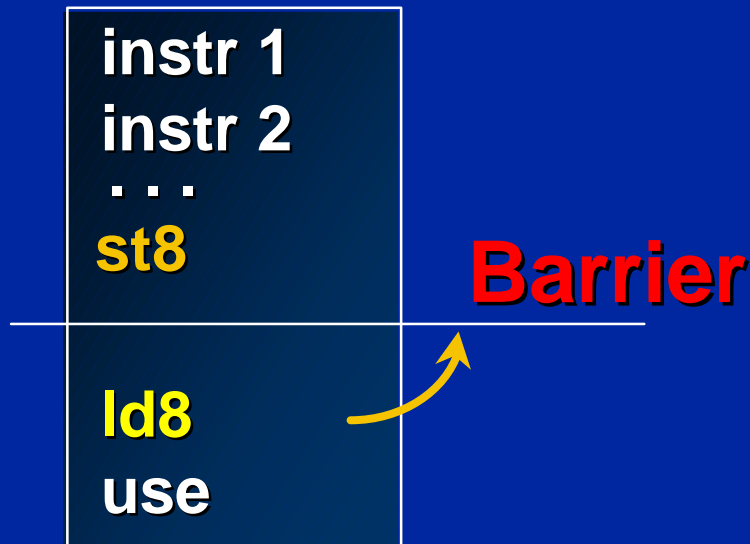
- The uses of speculative data can also be executed speculatively
- Provides additional scheduling flexibility to achieve greater parallelism

New!

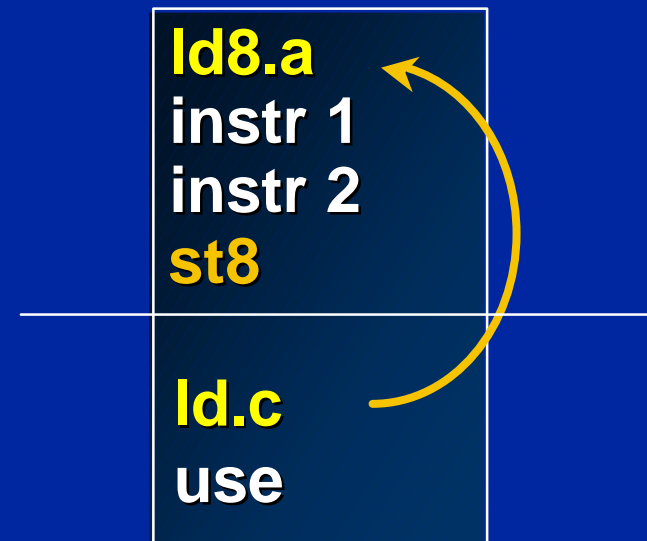
Introducing Data Speculation

- Compiler can issue a load prior to a preceding, possibly-conflicting store

Traditional Architectures



IA-64



Unique feature to IA-64





Exception Handling

Traditional Method 1

```
instr 1  
instr 2  
Load  
...  
br
```

use

Only when you
know it's safe

Traditional Method 2

```
instr 1  
instr 2  
Load  
...  
br
```

```
Load  
compare  
...  
use
```

Extra baggage
to ensure integrity

IA-64

```
ld.s  
instr 1  
instr 2  
br
```

```
chk.s  
use  
Home Block
```

;Exception Detection

NaT Bit

;Exception Delivery

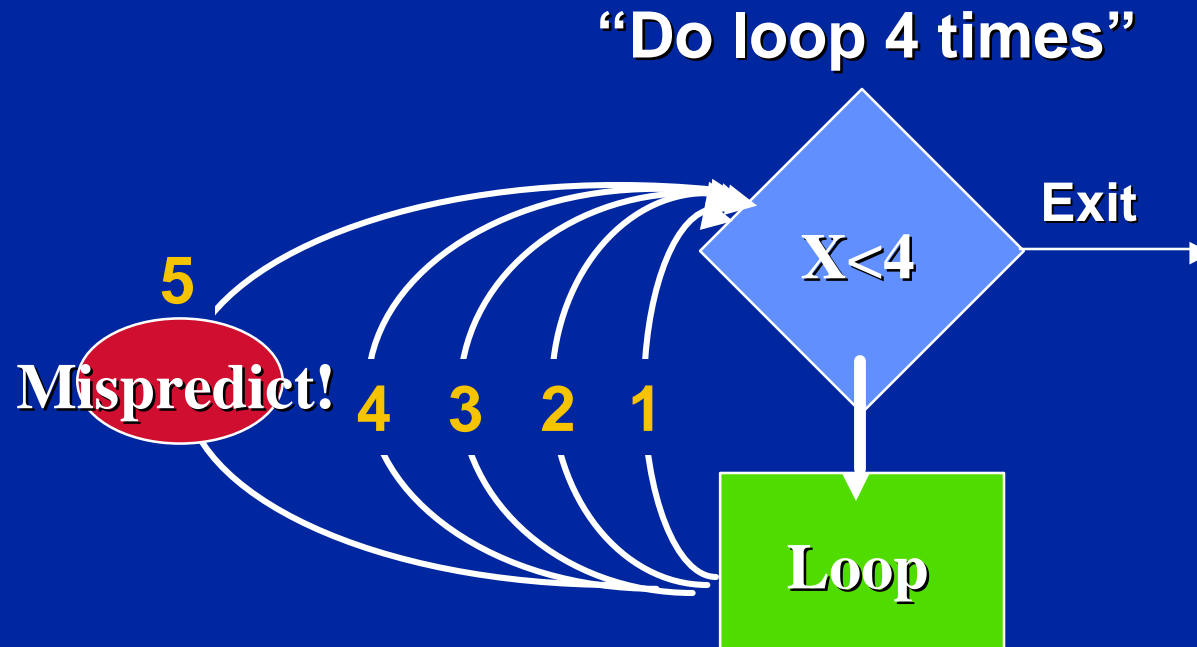
NaT bit ensures
integrity of data
without penalty

Speculation Benefits

- Reduces impact of memory latency
 - ◆ Performance improvement at 79% when combined with predication*
- Greatest improvement to code with many cache accesses
 - ◆ Large databases
 - ◆ Operating systems
- Scheduling flexibility enables new levels of performance headroom

New!

Loop Handling



IA-64 has special register to avoid #5:
called Loop Counter (LC)



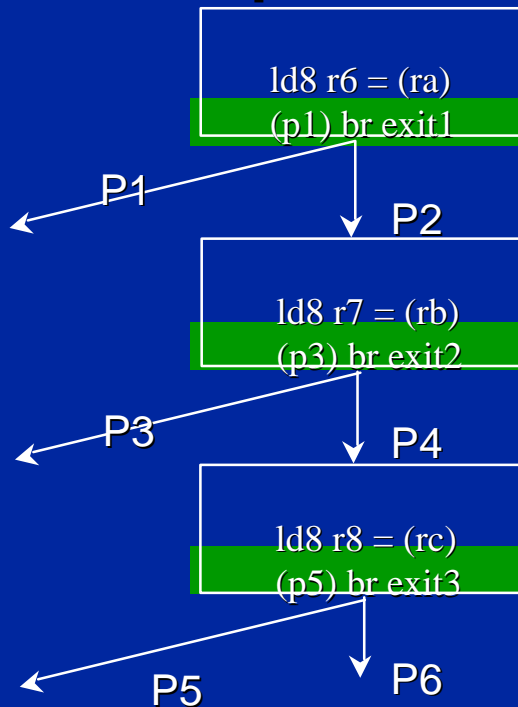
Improves integer code performance



New!

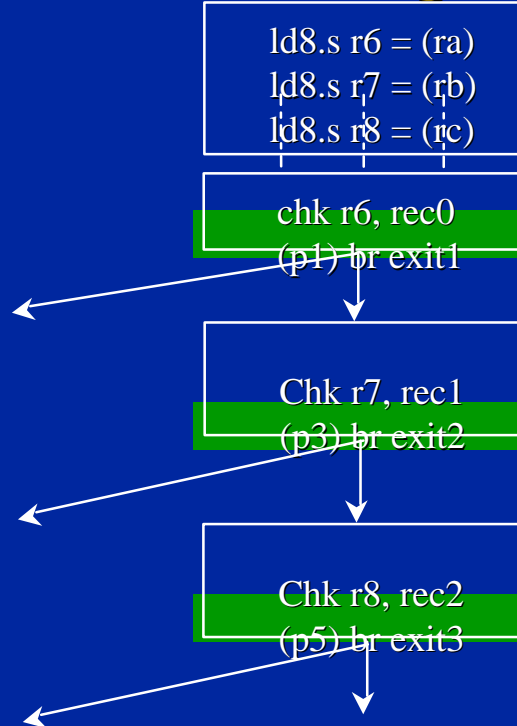
Multi-way Branch

w/o Speculation

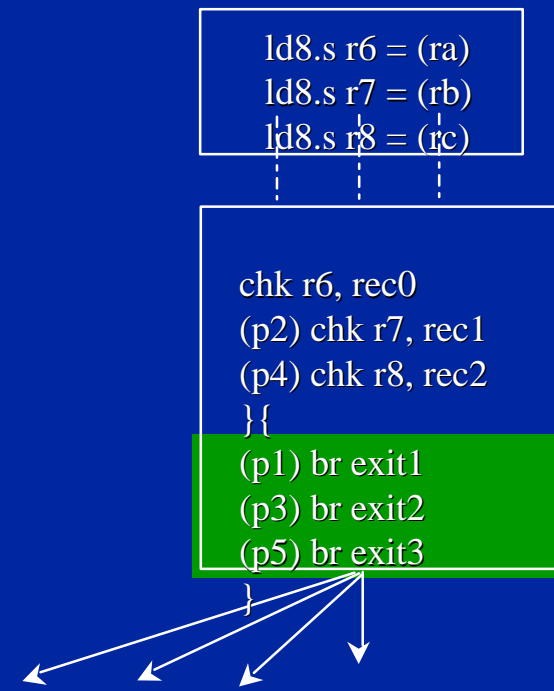


3 branch cycles

Hoisting Loads



IA-64



1 branch cycle

- Multiway branches: more than 1 branch in a single cycle
- ◆ Chaining multiway branches allows n-way branching

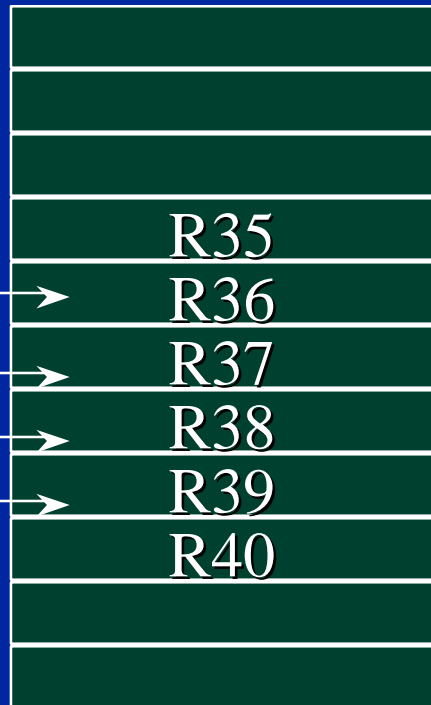
New!

Introducing Rotating Registers

Traditional Arch

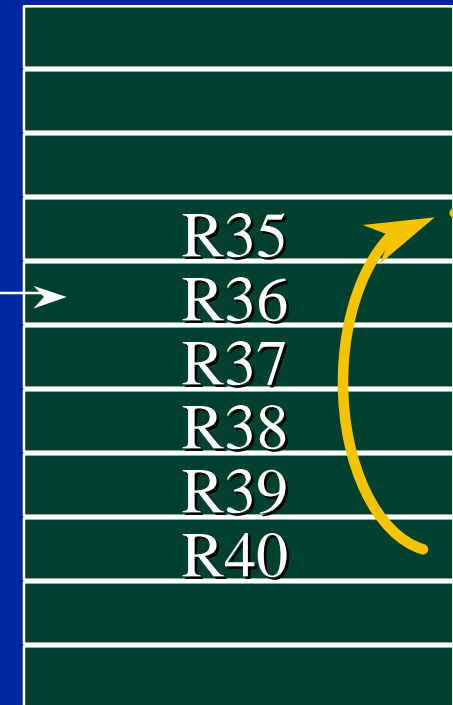
IA-64

load in R36
load in R37
load in R38
load in R39



4 instructions

load in R36



1 instruction

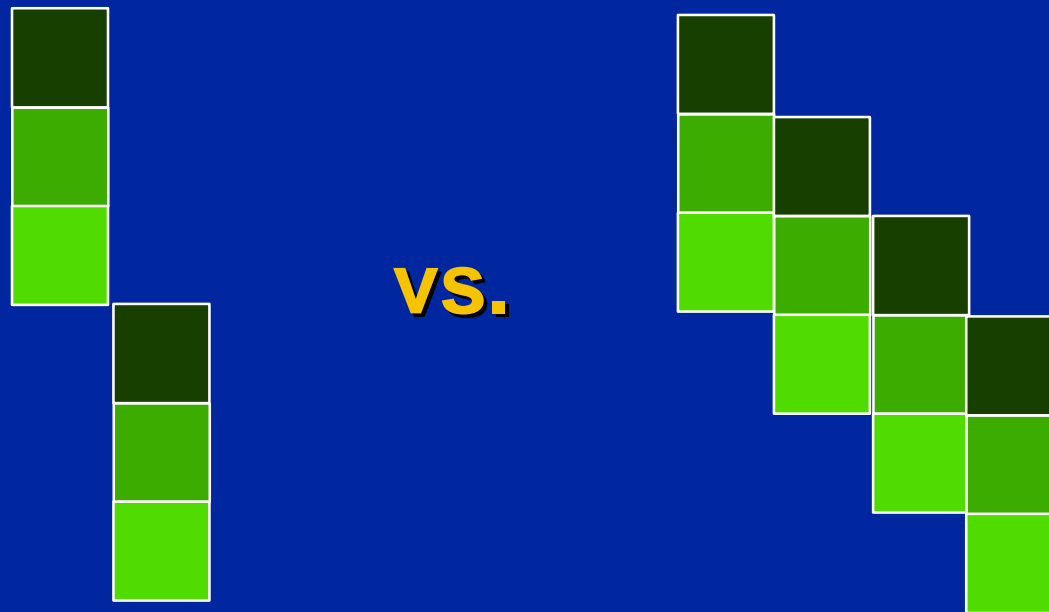
IA-64 Registers rotate, load stays the same



New!

Software Pipelining

- Overlapping execution of different loop iterations



- More iterations in same amount of time

IA-64 Offers a Unique Approach

Software Pipelining Benefits

- **Loop pipelining maximizes performance; minimizes overhead**
 - ◆ Avoids code expansion of unrolling and code explosion of prologue and epilogue
 - ◆ Smaller code means fewer cache misses
 - ◆ Greater performance improvements in higher latency conditions
- **Reduced overhead allows S/W pipelining of small loops with unknown trip counts**
 - ◆ Typical of integer scalar codes

Reviewing What's New:

- Parallel compares
- Tbit
- Nat bits
- Deferral
- Hoisting uses
- Propagation
- Branch instructions
- Static prediction
- Loop branches
- LC register
- EC register
- Multiway branch
- Branch registers
- Register rotation
- Predicate rotation
- RRBs

Feature Comparison

Traditional Architectures	IA-64
Dynamic branch prediction	Static Prediction and Predication enhance dynamic prediction to reduce mispredict penalties
Conditional moves limited in applicability and require additional instructions	Predication widely applicable, parallel compares further enhance benefit
Non faulting loads limited to certain conditions or require additional instructions	Control and data speculation enable greater scheduling freedom of loads
Software pipelining limited to large loops due to code size explosion	Rotating registers and rotating predicates allow wide application of software pipelining performance benefits without the code growth

Summary

- **Speculation reduces memory latency**
 - ◆ IA-64 removes recovery from critical path
 - ◆ Good for variety of server applications
- **Predication removes branches**
 - ◆ Logical and / or of compares increases parallelism
 - ◆ Good for large databases applications
- **S/W pipelining support with minimal overhead enables broad usage**
 - ◆ Performance for small integer loops with unknown trip counts as well as monster FP loops